On Local Algorithms for Topology Control and Routing in Ad Hoc Networks

Lujun Jia and Rajmohan Rajaraman
College of Computer Science
161 Cullinane Hall
Northeastern University
Boston, MA 02115, USA
{lujunjia,rraj}@ccs.neu.edu

Christian Scheideler
Department of Computer Science
Johns Hopkins University
3400 N. Charles Street
Baltimore, MD 21218, USA
scheideler@cs.jhu.edu

Abstract

An ad hoc network is a collection of wireless mobile hosts forming a temporary network without the aid of any fixed infrastructure. Indeed, an important task of an ad hoc network is to determine an appropriate topology over which high-level routing protocols are implemented. Furthermore, since the underlying topology may change with time, we need to design routing algorithms that effectively react to dynamically changing network conditions.

This paper studies algorithms for topology control and routing in ad hoc networks. We analyze the performance of the algorithms under three measures, *throughput*, which is the rate at which packets can be delivered for arbitrary communication patterns, *space* overhead, and the total *energy* consumed due to packet transmissions. Energy consumption is an important performance measure for ad hoc networks since the battery power of mobile nodes is usually limited.

We show that for any collection of nodes in the 2-dimensional Euclidean plane, a simple local algorithm identifies a connected constant degree graph that contains energy-efficient paths between every pair of nodes. We present routing algorithms designed for a model in which the topology as well as packet injections are under adversarial control. We show that a simple routing algorithm based on a local balancing approach achieves a constant-factor approximation with respect to both throughput and energy, when compared with any routing schedule. We also extend our algorithms and analyses for both topology control and routing to account for transmission interference, another important performance-limiting aspect of wireless communication.

Keywords

Distributed algorithms - Mobile computing and communication - Ad hoc networks - Routing - Spanners - Adversarial model - Competitive ratio

1 Introduction

An ad hoc wireless network consists of a collection of geographically dispersed nodes communicating with one another over a wireless medium using paths that may traverse multiple nodes. An ad hoc network differs from both wired and cellular networks in that there is no wired infrastructure and the communication capabilities of the network are usually constrained by the limited battery power of the nodes. While primary applications of ad hoc networks are in the military domain [18], the rapid advent of mobile telephony and a plethora of personal digital assistants has brought to the fore a number of potential commercial applications of ad hoc networks. Examples are disaster relief, conferencing, home networking, sensor networks, personal area networks, and embedded computing applications [33].

The absence of a fixed infrastructure in ad hoc networks implies that an ad-hoc network does not have an associated fixed topology. Hence, the nodes themselves have to form a connected topology to enable communication among them. There are several factors that influence the topology of an ad hoc network. Some are controllable such as the transmission power of individual nodes and antenna direction, while others are uncontrollable such as node mobility, weather, and noise. Furthermore, since wireless nodes transmit by broadcasting within a certain (potentially variable) transmission range, two different simultaneous transmissions may *interfere*, and neither may succeed. For a given topology, we also need to identify routes and schedule packet movements so to ensure high throughput and to minimize energy consumption, an important measure for communication in ad hoc networks.

In general, designing optimal communication protocols in ad hoc networks is hard. For instance, it is known that finding a schedule for a set of packets in an ad hoc network of n nodes that completes in time within even an $O(n^{1-\varepsilon})$ factor of optimal, is NP-hard, for any constant $\varepsilon > 0$ [1]. The sheer complexity of establishing communication in ad hoc networks suggests a layered approach, addressing the following questions:

- **How to set up a topology that guarantees connectivity?** Distributed algorithms that address this problem will be called *topology control protocols*. A naive solution that is wasteful in both energy consumption and maintenance overhead is to simply connect each node to all other nodes within its maximum transmission range. In order to increase scalability and reduce interference, it is more desirable to maintain only a constant number of direct links for each node at any point of time, while trying to ensure that the topology offers energy-efficient routes between any pair of nodes. Note that just connecting each node to its closest k neighbors may provide energy-efficient routes but does *not* guarantee connectivity or a constant degree per node.

- **How to select connections provided by the topology to allow non-interfering transmissions of packets?** A topology that ensures connectivity necessarily contains edges that interfere with each other. Thus, a way has to be found to schedule the use of these edges. Algorithms for this problem will be called *medium access control (MAC) protocols*.

- **How to route packets along non-interfering connections?** Given an underlying topology, which may be dynamically changing, we need to determine routes for individual packets and decide which packet to schedule if several packets contend to use an edge at the same time. We will refer to algorithms for this problem as *routing protocols*.

1.1 Our results

In this paper, we consider the performance of simple local algorithms for topology control, medium access, and routing in ad hoc wireless networks. To the best of our knowledge, this is the first study in which all of these issues have been addressed and analyzed.

Our first result concerns a local algorithm for computing a *constant-degree, energy-efficient* topology for an arbitrary distribution of ad hoc network nodes in the 2-dimensional Euclidean plane. Let V be a set of nodes in the 2-dimensional plane. We adopt the following standard model for energy consumption. The energy consumed due to a direct transmission from u to v is given by $|uv|^\kappa$, where $\kappa \geq 2$ is a constant and $|uv|$ is the Euclidean distance between u and v. The preceding formula for energy consumption, which is discussed in more detail in Section 2.2, follows from a standard power attenuation model adopted for wireless transmissions [35, 41]. The total energy used for delivering a packet from source s to destination t along a path P is simply the sum of the energy used for all the edges in P. We define the *energy-stretch* of a path P between vertices u and v to be the ratio of the energy of P to the energy of the minimum-energy path between u and v. The energy-stretch is a variant of the well-known measure of distance-stretch, which for a path P is the ratio of the length of P to the the minimum distance between u and v.

- We show that a simple local-control algorithm, proposed by Li et al [32], identifies an $O(1)$-degree graph $\mathcal{N}$ on V such that for any two nodes u and v, there exists a path in $\mathcal{N}$ between u and v that has $O(1)$ energy-stretch. For the special case of *civilized graphs*, in which it is assumed that the ratio of the maximum edge length to the minimum edge length is bounded by a constant, the same algorithm achieves $O(1)$ distance-stretch for any two nodes u and v. Our result, which is presented in Section 2, is related to work done on proximity graphs in computational geometry and may be of independent interest.

A topology control algorithm provides an underlying network over which a suitable routing mechanism can be implemented. Since this network is computed online and may further change due to uncontrollable factors (as discussed above), we need to design routing algorithms that react to dynamically changing network conditions.

First, we consider a scenario in which a topology control protocol and a MAC protocol given that provides edges to the routing layer that can be used without interference. We investigate the performance of our routing algorithm under the situation that the MAC protocol and the packet injections show adversarial behavior. More precisely, we assume that there is an adversary that is allowed to inject an arbitrary number of packets and that can select an arbitrary set of non-interfering edges at any time step. We also associate a *cost* with each edge that represents, for example, the energy usage for transmission along the edge, and may change from one step to another.

- Under the above adversarial model, we present a routing algorithm in Section 3 that is based on a simple, local balancing approach. For any sequence of adversarial packet injections and edge activations and for any constant $\varepsilon > 0$, our algorithm successfully delivers a $1 - \varepsilon$ fraction of the packets at an average cost that is within an $O(1/\varepsilon)$ factor of optimal, assuming that the node buffer sizes in our algorithm are larger than the buffer sizes used in an optimal schedule by a factor of essentially $O(\bar{L}/\epsilon)$, where $\bar{L}$ is the average path length used for successful packets in an optimal solution. While algorithms based on local balancing have been extensively studied before, this is the first study that models transmission costs; it is somewhat surprising that a local-control algorithm achieves a constant-factor approximation with respect to *both throughput and average cost*, when compared with any other routing schedule. We also note that the generality of our adversarial model implies the applicability of our result in diverse scenarios involving dynamic networks.

An important assumption in the above result is that transmissions across all of the edges in the network can be scheduled simultaneously. As mentioned at the outset, wireless nodes transmit by broadcasting and, therefore, transmissions are prone to interference, even when the nodes are able to adjust their transmission ranges. We adopt a standard model for interference, that is described in Section 2.4. Our next set of results address the impact of interference on the throughput achievable on the topology $\mathcal{N}$ and on the throughput achieved by our local balancing algorithm.

- We show that for any communication pattern, the local balancing algorithm, when applied to network $\mathcal{N}$ using a simple randomized symmetry-breaking technique for resolving interference, achieves throughput within $\Omega(1/I)$ of the optimal achievable on *any topology*, where I is the maximum number of edges that any edge in $\mathcal{N}$ interferes with. If the n nodes are distributed uniformly at random in the plane, then we show that $I = O(\log n)$ whp[1], thus implying that our local algorithms achieve a throughput within an $O(\log n)$ factor of any other routing algorithm on any topology. These results follow from our analyses in Sections 2.4 and 3.3.

- Finally, we also show in Section 3.4 that for the special case where the transmission range of every node is uniform and *fixed*, one can achieve expected throughput which is optimal to within constant factors.

1.2 Related work

The topology control algorithm that we analyze in this paper was first proposed by Li et al [32] and is a variant of a graph introduced by Yao [44] for connecting nodes in Euclidean space. In the Yao graph, which is also commonly referred to as the θ-graph, one partitions the space around each node into *sectors* of a fixed angle and connects the node to the nearest neighbor in each sector. It can be easily shown that the Yao graph contains paths of $O(1)$ energy-stretch connecting any two nodes. In fact, the Yao graph satisfies the stronger property of being a *spanner*; that is, for any two nodes u and v, the Yao graph contains a path connecting u and v, the length of which is within a constant factor of the Euclidean distance between u and v. (Note that a spanner always has constant energy-stretch.) The maximum degree of the Yao graph is $\Omega(n)$ in the worst case, however. One can obtain a bounded-degree subgraph of the Yao graph that is also a spanner by processing the edges in order by length and adding an edge (u, v) to the subgraph if there is no other edge (u, w) or (v, w) already added and having an angle close to that of (u, v) [36] (a related idea is used in [6]. A topology control algorithm due to Wattenhoffer et al [43] (also see [31]) adopts a similar approach to convert the Yao graph to a constant-degree spanner. All of the suggested approaches, however, rely on a global ranking of the edges and it is not apparent how to implement such a postprocessing of the Yao graph edges without network-wide communication. Our topology control result shows that there exists a simple local postprocessing of the Yao graph that maintains the $O(1)$ energy-stretch property while bringing the maximum degree down to a constant. We also analyze the throughput-efficiency of the resultant topology for arbitrary and random node distributions. Variants of the Yao graph are also studied in [23] under a civilized graph model of node locations, which assumes that the ratio of the maximum edge length to the minimum edge length is bounded by a constant.

The Yao graph and the variants discussed above are closely related to a class of graphs referred to as *proximity graphs*, in which the graph edges are determined by the proximity among the nodes in Euclidean space. Proximity graphs include relative neighborhood graphs and Gabriel graphs [39]. While the relative neighborhood graph has polynomial energy-stretch, a Gabriel graph, by definition, has shortest paths with respect to the ℓ_2-norm and hence has optimal energy paths. The Gabriel graph, however, has $\Omega(n)$ degree in the worst case. Another geometric structure that leads to a spanner is the Delaunay triangulation of the set of points; without additional restrictions, however, the Delaunay triangulation graph may include edges much longer than the transmission range of a node. It has been shown that restricted Delaunay graphs [21], in which we only include Delaunay edges with a limited fixed transmission radius, are also spanners. The maximum degree of restricted Delaunay graphs is $\Omega(n)$ in the worst case, however. For a comprehensive survey on geometric spanners and other structures in geometric network design, see [17].

In recent years, a number of routing protocols have been proposed for ad hoc networks. A recent survey may be found in [38]. Most of these protocols rely on heuristics and, as such, do not provide provable worst-case guarantees. Our work is also related to routing protocols that exploit the underlying geometry of the network [25, 29].

Our results on adversarial routing build on a series of studies in adversarial queuing theory, which was initiated by Borodin et al. [14]. Other work on adversarial queuing theory includes [5, 19, 20, 22, 37, 40]. In these studies it is assumed that the adversary has to provide a path for every injected packet and reveals these paths to the system. The paths have to be selected in a way that they do not overload the system. Hence, it only

[1]We use the abbreviation "whp" throughout the paper to mean "with high probability" or, more precisely, "with probability $1 - n^{-c}$, where n is the number of nodes in the network and c is a constant that can be set arbitrarily large by appropriately adjusting other constants defined within the relevant context."

remains to find the right queueing discipline (such as furthest-to-go) to ensure that all of the packets can reach their destination.

In the context of *packet routing algorithms*, the study of adversarial models was initiated by Awerbuch, Mansour and Shavit [12] and further refined by [4, 8, 10, 11, 20]. In the model adopted by these studies, the adversary *does not* reveal the paths to the system, and therefore the routing protocol has to figure out paths for the packets by itself. Based on work by Awerbuch and Leighton [11], Aiello et al. [4] show that there is a simple distributed routing protocol that keeps the number of packets in transit bounded in a dynamic network if, roughly speaking, in each window of time the paths selected for the injected packets require a capacity that is below what the available network capacities can handle in the same window of time.

Awerbuch et al. [8] study the problem of sending packets to a single destination in a dynamic network, using an adversarial model in which the adversary is allowed to control the network topology and packet injections as it likes, as long as for every injected packet it can provide a schedule to reach its destination. They show that even for the case that the network capacity is fully exploited, the number of packets in transit is bounded at any time. Recently, Awerbuch et al. [9] extended these results to arbitrary anycasting and multicasting situations and showed that simple balancing strategies achieve a throughput that can be brought arbitrarily close to a best possible throughput. Our work generalizes the results of [9] to incorporate edge costs. We also augment the algorithm to account for interference. We note that all of the above work in the adversarial routing area, including this current paper, is based on simple load balancing schemes first described in [12], and refined in [2, 3, 4, 8, 9, 10, 11] for various routing purposes.

2 Topology control

We consider a set V of n nodes in a 2-dimensional plane, in which each node can directly communicate with every node within a maximum distance D. Let $G^* = (V, E)$ denote the *transmission graph* that contains an edge between two nodes u and v if they can directly communicate with each other. We assume throughout this paper that G^* is connected. For each edge (u, v) in E, we associate a *cost* $c(u, v) = |uv|^\kappa$, for $\kappa \geq 2$. The cost of a path P between u and v is the sum of the cost of the edges along P.

We now elaborate on the assumptions made in our model. The parameter D represents the maximum transmission range of any node. The cost assigned to an edge represents the transmission energy expended and is based on a standard power attenuation model [35], in which the receiving power at any receiver is given by $\Theta(\frac{P_T}{d^\kappa})$, where P_T is the transmission power, d is the distance between the transmitter and the receiver and $2 \leq \kappa \leq 4$ is a constant. Thus, if we assume that each node has the same power reception threshold for signal detection, then the transmission power, and hence the energy, required for transmission over edge (u, v) is proportional to $|uv|^\kappa$, which is what we assign as energy cost for edge (u, v). We assume that each node is able to adjust its transmission power according to the distance to its receiver [43, 34], as long as the power does not exceed the maximum power needed to transmit to a distance of D.

Given an arbitrary collection of nodes forming a transmission graph G^*, we seek a distributed algorithm that identifies a low-degree subgraph of G^* that contains energy-efficient paths and admits high throughput. We capture the energy-efficiency of a subgraph H by its energy-stretch, which we now define. For any subgraph H of G^* and any nodes u and v, define $E^H_{u,v}$ to be the cost of the path with least cost in H. We define the *energy-stretch* of a subgraph H to be maximum ratio, over all nodes u and v, of $E^H_{u,v}$ to $E^{G^*}_{u,v}$. The main result of this section is that for any distribution of the n nodes, a simple local algorithm computes an $O(1)$-degree topology $\mathcal{N}$ with $O(1)$ energy-stretch. We note that the results of Wang et al. [42] establish the constant energy-stretch property of $\mathcal{N}$ for the special case of civilized graphs [27]. For this special case, we show in this section that the topology $\mathcal{N}$ actually achieves $O(1)$ distance-stretch for any two nodes. Thus, the constant energy-stretch property follows directly from the constant distance-stretch property. For a general distribution of nodes, however, we have not been able to resolve whether $\mathcal{N}$ is a spanner and we leave this question as an open problem at this time. The algorithm is described in Section 2.1, and the analysis of energy-stretch and distance-stretch are presented in Section 2.2 and Section 2.3 respectively.

We also evaluate the topology $\mathcal{N}$ on the basis of the throughput achievable for arbitrary communication patterns. The throughput achievable on a topology depends on the degree to which the edges of the topology interfere. A formal definition of the interference model and the analysis of throughput are give in Section 2.4.

2.1 Algorithm

In this section, we describe the topology control algorithm proposed in [32]. The algorithm is parametrized by an angle $\theta \leq \pi/3$. We refer to the algorithm as ΘALG. Each node $u \in V$ divides the $360°$ space into $2\pi/\theta$ *sectors*. For any two nodes u and v, we let $S(u,v)$ denote the sector of u containing node v. In the following description, we assume, without loss of generality, that all pairwise distances among the n nodes are unique. (If the distances are not unique, then a simple tie breaking scheme can be used to enforce the assumption.) The ΘALG determines a subgraph $\mathcal{N} = (V, E)$ of G in two phases:

1. Each node u computes $N(u)$ which consists of all nodes v such that v is the node nearest to u in $S(u,v)$.

2. Edge $(u,v) \in E$ is in E if v is the nearest node in $S(u,v)$ such that $u \in N(v)$ or u is the nearest node in $S(v,u)$ such that $v \in N(u)$.

Let $\mathcal{N}_1$ denote the graph obtained after the first phase of the algorithm; that is $\mathcal{N}_1 = (V, E_1)$ where $(u,v) \in E_1$ if $u \in N(v)$ or $v \in N(u)$. The graph $\mathcal{N}_1$ is identical to the Yao graph with θ-degree sectors. One can easily prove by an induction on pairwise distances that the distance between two nodes u and v in the graph is $O(|uv|)$, and hence that $\mathcal{N}_1$ is a spanner. It follows then that $\mathcal{N}_1$ also has $O(1)$ energy stretch.

While the total number of edges in $\mathcal{N}_1$ is $O(n)$, the degree of a node may be as large as $\Omega(n)$ in the worst-case. One can construct a constant-degree subgraph of $\mathcal{N}_1$ that preserves the spanner property by processing the edges in order of decreasing length, and eliminating edges that do not decrease the distance between endpoints by more than a constant-factor [43]. Such a postprocessing step, however, takes communication time proportional to the diameter of the network. Instead, the second phase of the algorithm above proposes a simple local step to eliminate certain edges from $\mathcal{N}_1$ so that the degree of each node is a constant. In Section 2.2, we will show that the resultant graph $\mathcal{N}$ has $O(1)$ energy-stretch.

Before going on to the analysis, we note that ΘALG can be implemented by three rounds of local message broadcasting and computation. In the first round, each node broadcasts a *Position* message containing its position, at maximum power P. After receiving the position information, each node u computes $N(u)$. In the second round, each node u broadcasts a *neighborhood* message containing $N(u)$ to $N(u)$. In the third round, each node u sends a *connection* message to the nearest node v (if any) in each sector such that u is in $N(v)$. The topology $\mathcal{N}$ has an edge (u,v) for any pair of nodes u and v that have exchanged a connection message.

It can be easily shown that the topology $\mathcal{N}$ is connected and has constant degree [42].

Lemma 2.1 ([42]) *$\mathcal{N}$ is connected and the degree of each node is at most $4\pi/\theta$.*

Proof: Assume for the purpose of contradiction that $\mathcal{N}$ is not strongly connected. Then, there exists at least one pair of nodes u, v, such that (u,v) is an edge in G^*, while there is no path from u to v in $\mathcal{N}$. Assume u, v is the pair with the shortest distance among those pairs that has no path in $\mathcal{N}$. We distinguish the following two cases:

Case 1: u is the nearest neighbor of v in $S(v,u)$, as shown in Figure 1. Since v adds u to $N(v)$ and (u,v) is not an edge in $\mathcal{N}$, there exists a node w in $S(u,v)$, such that $u \in N(w)$ and $(u,w) \in E$. In this case, $|uw| < |uv|$ and $|wv| < |uv|$, because $\angle wuv \leq \theta \leq \pi/3$.

Case 2: u is not the nearest neighbor of v in $S(v,u)$, as shown in Figure 2. In this case, let w be v's nearest neighbor in $S(v,u)$. Thus, we have $|wv| < |uv|$ and $|uw| < |uv|$, since $\angle wvu \leq \theta \leq \pi/3$.

Since in both cases, there is no path from u to v in $\mathcal{N}$, there is either no path from u to w or no path from w to v. This contradicts the assumption that u, v is pair of nodes with the shortest distance among those pairs without a path in $\mathcal{N}$.

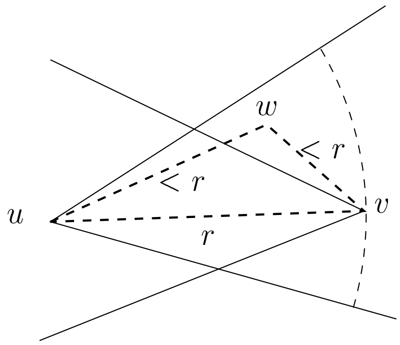

Figure 1: u is v's nearest neighbor.

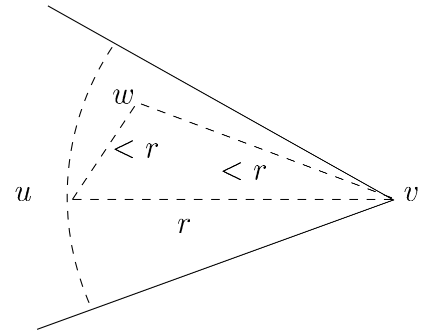

Figure 2: u is not v's nearest neighbor.

We now consider the degree of any node. Node u has an edge (u, v) only if either v is the nearest node in $N(u)$ or is the nearest node in the sector $S(u, v)$ such that $u \in N(v)$. Thus, each node has at most two edges in each sector. Summing up over the $2\pi/\theta$ sectors, we have the desired bound. ■

2.2 Analysis of energy-stretch

In this section, we show that $\mathcal{N}$ has $O(1)$ energy-stretch. Let $E_{u,v}$ denote the cost of the minimum-cost path from u to v in $\mathcal{N} = (V, E)$. We will show that for any pair of nodes u, v, $E_{u,v}$ is within a constant factor of the minimum cost to transmit from u to v in G^*. Since, the transmission along any edge (u, v) in G^* incurs cost $|uv|^\kappa$, it suffices to show that for any edge $(u, v) \in G^*$, $E_{u,v}$ is $O(|uv|^\kappa)$. Our main theorem is as follows

Theorem 1 *For θ sufficiently small, $E_{u,v} = O(|uv|^\kappa)$, for any edge (u, v) in G^*.*

The proof of Theorem 1 proceeds by induction on pairwise distances. A challenge in establishing Theorem 1 is that unlike proximity graphs such as the Yao graph [44], Gabriel graph and some of its variants (such as β-skeletons with $\beta < 1$) [17], the minimum-cost path in $\mathcal{N}$ from a node u to another node v may traverse nodes that are farther from v than u is. We are able to overcome this hurdle by sufficiently characterizing such a path so as to place an upper bound on the cost.

For our proof of Theorem 1, we need a series of technical lemmas that establish relationships among node distances and relative orientation. These lemmas and their proofs are presented in Section 2.2.1. The proof of Theorem 1 is given in Section 2.2.2.

2.2.1 Technical lemmas

Lemma 2.2 *For any $\triangle ABC$ with $|AC| \leq |BC|$ and $\angle ACB \leq \pi/3$, $c|AB|^2 + |AC|^2 \leq c|BC|^2$ for $c \geq \frac{1}{2\cos(\angle ACB) - 1}$.*

Proof: Without loss of generality, let $|BC| = 1, |AC| = x$ and $\angle ACB = \alpha$. Then, we have,

$$c|AB|^2 + |AC|^2 - c|BC|^2 = c(1 + x^2 - 2x\cos\alpha) + x^2 - c = (1 + c)x^2 - 2cx\cos\alpha \leq 0,$$

for $c \geq \frac{1}{2\cos\alpha - 1}$. ■

Lemma 2.3 *For any $\triangle ABC$ with $|BC| \leq |AC| \leq |AB|$ and $\angle BAC \leq \pi/6$, $|BC| \leq \frac{|AB|}{2\cos\angle BAC}$.*

Proof: Let $|AB| = 1$ and $\angle BAC = \alpha$. Then, we have $|BC| \leq \max\{2\sin\frac{\alpha}{2}, \frac{1}{2\cos\alpha}\} = \frac{1}{2\cos\alpha}$, for $\alpha \leq \pi/6$. ■

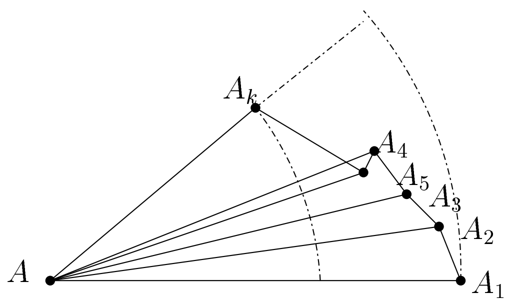

Figure 3: Node sequence A_i.

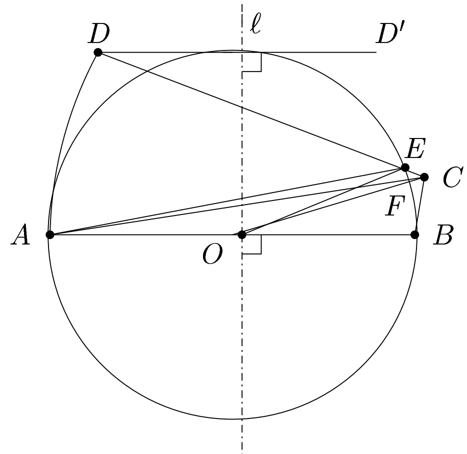

Figure 4: $|EC|$ is smaller than $|BC|$.

Lemma 2.4 *Let $A, A_1, A_2, \dots, A_k$ be a set of points, such that $|AA_i| \geq |AA_{i+1}|$ and $0 \leq \angle A_i A A_{i+1} \leq \theta$. If $\angle A_1 A A_k = \alpha$, then $\sum_{i=1}^{k-1} |A_i A_{i+1}|^2 \leq (|AA_1| - |AA_k|)^2 + 2|AA_1|^2 \frac{\alpha}{\theta}(1-\cos\theta)$.*

Proof: Consider Figure 3. Let $|AA_i| = d_i$ and $\angle A_i A A_{i+1} = \alpha_i$. Since $\cos x$ is a concave function on $[0, \pi/2]$ and $0 \leq \alpha_i \leq \theta$, we have

$$\sum_{i=1}^{k-1} \cos\alpha_i \quad \geq \quad k - 1 - \frac{\alpha}{\theta} + \frac{\alpha}{\theta}\cos\theta.$$

Thus, we have

$$\begin{aligned} \sum_{i=1}^{k-1} |A_i A_{i+1}|^2 &= \sum_{i=1}^{k-1} (d_i^2 + d_{i+1}^2 - 2 d_i d_{i+1} \cos\alpha_i) \\ &\leq (d_1 - d_k)^2 + 2d_1^2 \sum_{i=1}^{k-1} (1 - \cos\alpha_i) \\ &\leq (d_1 - d_k)^2 + 2d_1^2 \frac{\alpha}{\theta}(1 - \cos\theta), \end{aligned}$$

which completes the proof. ∎

Lemma 2.5 *Let A and B be any two points, and O be the center of line segment (A, B). Let D be the point such that $|BD| = |AB|$ and $\angle DBA = \pi/6$. Let C be a point outside $C(O, |OA|)$ such that $|AC| \leq |AB|$ and $\angle CAB < \pi/12$ and C, D are on the same side of (A, B). Let E be the intersection of (C, D) with circle $C(O, |OA|)$. We have $\angle EAB \leq 2\angle CAB$.*

Proof: Consider Figure 4. Let F be the intersection of (C, O) with $C(O, |OA|)$. Let $\angle CAB = \alpha$. Since $|OC| > |OA|$, we have $\angle BOC \leq 2\alpha$. Thus,

$$\angle BCO = \pi - \angle BOC - \angle CBO \geq \pi - 2\alpha - \pi/2 = \pi/2 - 2\alpha,$$

since $\angle CBO < \pi/2$. We also have

$$\angle OCE = \angle ACD + \angle ACO = \angle CAB + \angle CDD' + \angle ACO < \pi/6 + 2\alpha \leq \pi/2 - 2\alpha,$$

since $\alpha \leq \pi/12$ and $\angle CDD' \leq \angle BDD'$. Thus, we have $\angle OCE \leq \angle OCB$. Thus, $|EF| > |FB|$, and $\angle CAE \leq \angle BAC$. This completes the proof. ∎

2.2.2 Proof of Theorem 1

Proof of Theorem 1: We prove that $E_{u,v} \le c|uv|^\kappa$ for an appropriately chosen c, for any pair of nodes $u, v \in V$. For simplicity, we assume $\kappa = 2$ in the following analysis, and it is easy to see that the results applies to any $\kappa > 2$. We prove this theorem by an induction on $|uv|$ for any $u, v \in V$.

Consider the base case where $|uv|$ is the minimum among any pair of nodes in V. This is trivial, since it is clear that $(u, v) \in E$ and $E_{u,v} = |uv|^2 \le c|uv|^2$, for $c \ge 1$. Consider nodes u, v. Assume for the purpose of induction that for any pair of nodes x, y with $|xy| < |uv|$, $E_{x,y} \le c|xy|^2$. We show in the following that $E_{u,v} \le c|uv|^2$. We distinguish between the following four cases:

Case 1**:** $v \in N(u)$, or $u \in N(v)$. In this case, u selects v as its nearest neighbor, so there exists node w ($w = v$ or $w \ne v$) in sector $S(v, u)$ with $\angle uvw \le \theta$ such that $|wv| \le |uv|$ and $(w, v) \in E$. Invoking Lemma 2.2, we have $E_{u,v} \le E_{u,w} + |wv|^2 \le c|uw|^2 + |wv|^2 \le c|uv|^2$, for $c \ge \frac{1}{2\cos\theta - 1}$.

Case 2**:** $v \notin N(u)$ and $u \notin N(v)$. In this case, if there exists a node w with $\angle uwv \ge \pi/2$, we have $E_{u,v} \le E_{u,w} + E_{w,v} \le c|uw|^2 + c|wv|^2 \le c|uv|^2$.

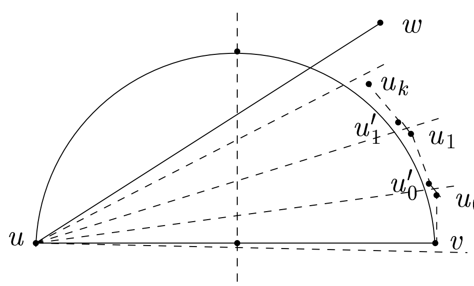

Figure 5: Sectors of node u.

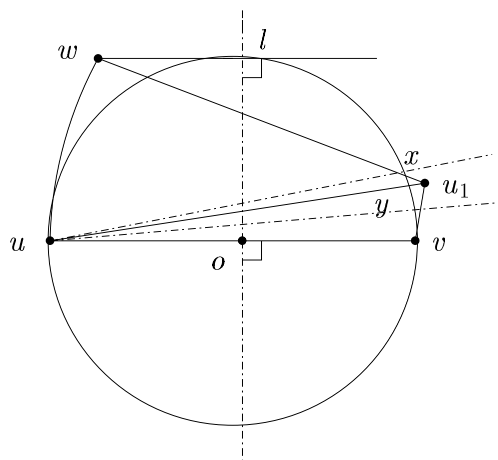

Figure 6: Selection of node z.

We now assume that $\angle uwv < \pi/2$ for any $w \in V$. Consider Figure 5. Let $\angle wuv = \pi/6$. We number the consecutive sectors of u as sector $0, 1, \ldots, k$, starting from $S(u, v)$. The number k is selected to be the largest number such that sectors 0 through k do not contain (u, w). We denote the nearest neighbor of u in sector i as u_i, and let u'_i be the node such that $u_i \in N(u'_i)$ and $(u_i, u'_i) \in E$, if such a node exists. Note that whether the sectors are above or below line segment (u, v) depends on the position of u'_0, which will be made clear in the following analysis. We the let u_t be the first node in the sequence $u_0, u_1, \ldots, u_k$, such that $|uu'_t| < |u'_t u_t|$. We distinguish between the following two cases:

Case 2.1**:** u_t exists for $0 \le t \le k$. In this case, $|u_i u'_i| \le |uu'_i| \le |uu_i|$ for $0 \le i \le t - 1$. We then have $\angle u'_i uu_i \le \angle u'_i u_i u \le \theta$. Note that u'_i can not lie in any sector j for $j \le i$, since otherwise u'_i would be the nearest neighbor of u instead of u_k. Thus, u'_i can only fall in sector $i + 1$, which implies the existance of u_{i+1} and $|uu_{i+1}| \le |uu'_i|$. It is obvious that u_0 exists, we then have a sequence of nodes $v, u_1, u_2, \ldots, u_t$, such that their distance to u is decreasing and $\angle vuu_1, \angle u_i uu_{i+1} \le 2\theta$ for $1 \le i \le t - 1$. Now consider $\triangle uu'_t u_t$. Since $\angle u_t uv \le \pi/6$ and $\angle uu_t v < \pi/2$, it is clear that $\frac{\sqrt{3}}{2}|uv| \le |uu_t| \le |uv|$; invoking Lemma 2.3, we have $|uu'_t| \le \frac{|uu_t|}{2\cos\theta} \le \frac{|uv|}{2\cos\theta}$. Invoking Lemma 2.4 on the sequence $v, u_1, u_2, \ldots, u_t$, we have $E_{u_t,v} \le$

$c(|uv| - |uu_t|)^2 + 2c|uv|^2 \frac{\pi}{12\theta}(1 - \cos(2\theta)) \le c(1 - \frac{\sqrt{3}}{2})^2|uv|^2 + \frac{c\pi}{6\theta}(1 - \cos(2\theta))|uv|^2$. Thus, we have

$$\begin{aligned} E_{u,v} &\le E_{u,u'_t} + E_{u'_t,u_t} + E_{u_t,v} \\ &\le c\frac{|uv|^2}{4\cos^2\theta} + |uv|^2 + c(1 - \frac{\sqrt{3}}{2})^2|uv|^2 + \frac{c\pi}{6\theta}(1 - \cos(2\theta))|uv|^2 \\ &\le c|uv|^2, \end{aligned}$$

for $c \ge 2.3$ and $\theta \le \pi/12$. This completes the proof for this subcase.

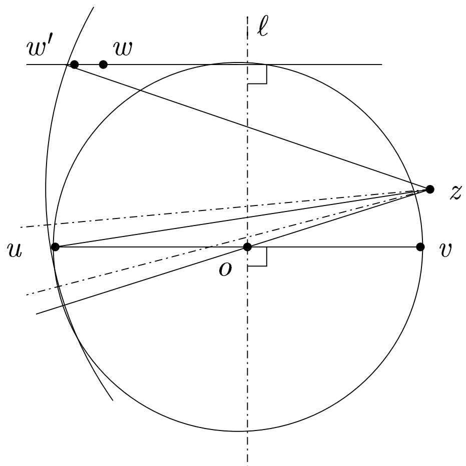

Figure 7: Sectors of node z.

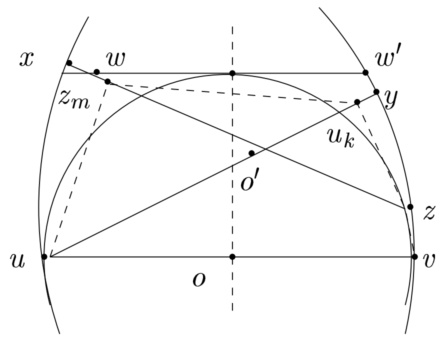

Figure 8: Induction without selecting edge E.

Case 2.2**:** u_t does not exist for $0 \le t \le k$. In this case, there exists a sequence of nodes $v, u_1, u_2 \ldots, u_k$, such that their distance to u is decreasing and $\angle vuu_1, \angle u_iuu_{i+1} \le 2\theta$ for $1 \le i \le k - 1$.

Consider Figure 6 and Figure 7. Let point o be the center of (u, v) and ℓ be the perpendicular bisector of line segment (u, v). We show in the following that a similar nodes sequence can be identified on the side of u. Let w, w' be two points such that $\angle wvu = \angle w'uv = \pi/6$ and $|wv| = |w'u| = |uv|$. Let x (resp., y) be the intersection of (u_1, w) (resp., (u_1, u)) with circle $C(o, |ou|)$. We now show that there exists a node z in $\triangle u_1xy$ such that

1. $\theta \le \angle zuv \le 4\theta$;
2. the nearest neighbor of z in sector $S(z, u)$ is on the same side of (u, v) as z;
3. any neighbor of z between ray (z, u) and (z, w) is on the same side of ℓ as u.

We select z to be the nearest node to u in $\triangle xu_1y$. Since u_1 is the nearest neighbor of u in sector 1, z is not in sector 1 of u (or $z = u$). Thus, $\angle zuv > \theta$. In the following analysis, we assume $\theta \le \pi/24$. Since $\angle u_1uv \le 2\theta < \pi/12$, invoking Lemma 2.5, we have $\angle zuv \le 2\angle u_1uv \le 4\theta$. It can be show that $\angle uzo \ge \angle zuv/2 \ge \theta/2$, which proves the second property. Since $\angle wzu \le \pi/6 + \angle zuv \le \pi/6 + 4\theta \le \pi/3$ and we choose z to be the nearest node to u in $\triangle xu_1y$, it is easy to show that the third property holds.

Now we are ready to establish the sequence of nodes in terms of z. Consider Figure 7. We number the consecutive sectors of z as $0, 1, \ldots, m$, starting from $S(z, u)$. The number m is selected to be the largest number such that sectors 0 through m do not contain (z, w). Denote the nearest neighbor of z in sector i as z_i, and let z'_i denote the node such that $z_i \in N(z'_i)$ and $(z_i, z'_i) \in E$, if such a node exists. We denote the intersection of ray (z, z_m) with (w, w') as x_m. We then put an additional constraint on m that $\angle z_mzu > 2\angle zuv$, in order to make the projection of x_m on (u, v) fall on the right side of u. It can easily be shown that $\pi/6 \le \angle wzu \le \pi/4$ for our z. Since $\angle zuv \le 4\theta$, $\theta \le \pi/60$ satisfies those requirement, and we assume this in the following analysis.

Let z_t be the first node in the sequence $z_0, z_1, \ldots, z_m$, such that $|z_t z'_t| > |z'_t z|$. We then distinguish the following two subcases:

Case 2.2.1**:** z_t exists for $0 \le t \le m$. Given $\theta \le \pi/60$, it is easy to give a bound on $|uz_t|$ that $|uz_t| \le \sqrt{2}|uv|/2$. It can be shown that $|z_t z'_t| \le |z_t z| \le |uv|$. Invoking Lemma 2.3, we have $|z'_t z| \le |z_t z|/(2\cos\theta) \le |uv|/(2\cos\theta)$. We also have $|zv| \le 2\sin(2\theta)|uv|$, since $\angle zuv \le 4\theta$. By induction, we then have

$$\begin{aligned}
E_{u,v} &\le E_{u,z_t} + E_{z_t,z'_t} + E_{z'_t,z} + E_{z,v} \\
&\le c|uz_t|^2 + |z_t z'_t|^2 + c|z'_t z|^2 + c|zv|^2 \\
&\le \frac{c|uv|^2}{2} + |uv|^2 + \frac{c|uv|^2}{4\cos^2\theta} + 4c\sin^2(2\theta)|uv|^2 \\
&\le c|uv|^2,
\end{aligned}$$

for $c \ge 4.9$ and $\theta \le \pi/60$.

Case 2.2.2**:** z_t does not exist for $0 \le t \le m$. In this case, we have a sequence of nodes $u, z_0, z_1, \ldots, z_n$, such that their distance to z is decreasing and $\angle uzz_0, \angle z_i z z_{i+1} \le 2\theta$. Consider Figure 8. Let w, w' be two points such that $\angle wvu = \angle w'uv = \pi/6$ and $|wv| = |w'u| = |uv|$. Let x, y be the intersections of $(u, u_k), (z, z_m)$ with $C(z, |zu|), C(u, |uv|)$ respectively. Note that $\angle xo'y$ is obtuse, we then have $|xz_n|^2 + |z_n u_k|^2 + |u_k y|^2 \le |xy|^2$. Point y is always below line (w, w'). From the definition of number k, $\angle yuw' \le 2\theta$, which implies that distance of y to line (w, w') is at most $2 * 2\theta|uy| \le 4\theta|uv|$. Now consider x. From the definition of m, we have $\angle xzw \le 2\theta$. If x is below line (w, w'), then the distance of x to line (w, w') is at most $4\theta|zx| \le 4\theta|uv|$. If x is above line (w, w'), then the distance is maximized when lines (z, w) and (z, x) are co-linear. Thus, the distance of x to line (w, w') is the product of $\sin(\angle zww')$ and $|zx| - |zw|$. Since $|zx| - |zw| = |zu| - |zw| \le |uv| - |zw| = |vw| - |zw| \le 4\theta$, and $\sin(\angle zww') \le \sin(\angle vww') = 1/2$. So, in either case, the distance of x to line (w, w') is at most 2θ. Thus, we have

$$\begin{aligned}
|xy|^2 &\le (\cos(\angle yuv)|uv|)^2 + (4\theta|zx| + 2\theta|uy|)^2 \\
&\le \cos^2(\frac{\pi}{6} - 2\theta)|uv|^2 + 36\theta^2|uv|^2,
\end{aligned}$$

since x, y may lie on different side of (w, w')

Recall that $\angle yuv \le \angle wuv = \pi/6$ and $\angle xzu \le \angle wzu \le \pi/4$. Invoking Lemma 2.4 on the node sequence $v, u_1, u_2, \ldots, u_k$ and sequence $u, z_0, z_1, \ldots, z_n$ respectively, we have

$$\begin{aligned}
E_{u,v} &\le E_{u,z_n} + E_{z_n,u_k} + E_{u_k,v} \\
&\le c(|zu| - |zz_n|)^2 + 2c|zu|^2\frac{\angle xzu}{2\theta}(1 - \cos(2\theta)) + c|z_n u_k| + c(|uv| - |uu_k|)^2 \\
&+ 2c|uv|^2\frac{\pi}{6}\frac{1 - \cos(2\theta)}{2\theta} \\
&\le c|xz_n|^2 + c\pi|uv|^2\frac{1 - \cos(2\theta)}{4\theta} + c|z_n u_k|^2 + c|u_k y|^2 + c\pi|uv|^2\frac{1 - \cos(2\theta)}{6\theta} \\
&\le c\cos^2(\frac{\pi}{6} - 2\theta)|uv|^2 + 36c\theta^2|uv|^2 + c\pi|uv|^2\frac{1 - \cos(2\theta)}{4\theta} + c\pi|uv|^2\frac{1 - \cos(2\theta)}{6\theta} \\
&\le c|uv|^2,
\end{aligned}$$

for any c and $\theta \le \pi/60$. This completes the proof of Theorem 1. ■

2.3 Analysis of distance-stretch for civilized graphs

A civilized graph, also called λ-precision graph, satisfies the following property: $\min\{\frac{|u_1 v_1|}{|u_2 v_2|}\} \ge \lambda$ for any nodes u_1, u_2, v_1, v_2 in the graph, and $0 < \lambda \le 1$ is a constant. This is a commonly used model for wireless ad hoc networks [42, 23], since wireless devices typically are not too close to each other. Let Θ-*distance* $D^{\mathcal{N}}(u, v)$ denote the minimum distance between u, v in topology $\mathcal{N}$. Our main theorem is as follows

Theorem 2 *If G^* is a civilized graph, then topology $\mathcal{N}$ has a distance-stretch of $O(1)$ for sufficiently small θ.*

To prove Theorem 2, we first note the spanner property of the Yao graph. Let *Yao-distance*, denoted by $D^{\mathrm{Y}}(u,v) = |u_0u_1| + |u_1u_2| + \ldots + |u_{k-1}u_k|$, be the length of the path from $u\ (= u_0)$ to $v\ (= v_0)$ in G^*, where u_i is the nearest neighbor of u_{i-1} in sector $S(u_{i-1}, v)$ for any $1 \leq i \leq k$. It is easy to see that each edge in this path is an edge in the Yao graph of G^*. It has been shown that $D^{\mathrm{Y}}(u,v) \leq \frac{1}{1-2\sin\frac{\theta}{2}}|uv|$ for any two nodes u and v [30]. We then show that the distance between any two nodes u, v in $\mathcal{N}$ is within a constant factor of their Yao-distance, given civilized graph. For $D^{\mathrm{N}}(u,v)$ in topology $\mathcal{N}$, we have the following

Lemma 2.6 *For any two nodes u, v in $\mathcal{N}$ defined on a civilized graph, $D^{\mathrm{N}}(u,v) = O(D^{\mathrm{Y}}(u,v))$ for sufficiently small θ.*

Proof: We invoke an induction on the pairwise Euclidean distance between any two nodes u, v in the civilized graph. Let c_1 denote the constant $\frac{1}{1-2\sin\frac{\theta}{2}}$. For the base case, (u,v) is the minimum length edge. We then have $D^{\mathrm{N}}(u,v) = D^{\mathrm{Y}}(u,v) \leq c_2 D^{\mathrm{Y}}(u,v)$ for any constant $c_2 \geq 1$, since (u,v) is an edge in both the Yao graph and topology $\mathcal{N}$, which is a subgraph of the Yao graph. Assume for the purpose of induction that for any nodes x, y with $|xy| < |uv|$, $D^{\mathrm{N}}(x,y) \leq c_2 D^{\mathrm{Y}}(x,y)$. We now consider nodes u, v, and distinguish the following three cases:

Case 1: Node v is not the nearest neighbor of u in sector $S(u,v)$. In this case, u has a nearest neighbor $w \neq v$ in sector $S(u,v)$. Since $\angle wuv \leq \theta \leq \frac{\pi}{3}$ and $|uw| < |uv|$, we have $D^{\mathrm{N}}(u,w) \leq c_2 D^{\mathrm{Y}}(u,w), D^{\mathrm{N}}(w,v) \leq c_2 D^{\mathrm{Y}}(w,v)$ by induction. Thus, we have

$$D^{\mathrm{N}}(u,v) \leq D^{\mathrm{N}}(u,w) + D^{\mathrm{N}}(w,v) \leq c_2 D^{\mathrm{Y}}(u,w) + c_2 D^{\mathrm{Y}}(w,v) = c_2 D^{\mathrm{Y}}(u,v).$$

The last equation is from the definition of Yao-distance, since w is the nearest neighbor of u in sector $S(u,v)$.

Case 2: Node v is the nearest neighbor of u in $S(u,v)$ and (u,v) is an edge in $\mathcal{N}$. This case is trivial since (u,v) is an edge in $\mathcal{N}$. Thus, $D^{\mathrm{N}}(u,v) = D^{\mathrm{Y}}(u,v) \leq c_2 D^{\mathrm{Y}}(u,v)$ for any constant $c_2 \geq 1$.

Case 3: Node v is the nearest neighbor of u in $S(u,v)$ but (u,v) is not an edge in $\mathcal{N}$. By the definition of ΘALG, there exists a node $w \neq u$ such that v is the nearest neighbor of w in sector $S(w,v)$ and (w,v) is an edge in $\mathcal{N}$. In this case, (w,v) is an edge in $\mathcal{N}$ and $D^{\mathrm{N}}(u,w) \leq c_2 D^{\mathrm{Y}}(u,w)$ since $\angle uvw \leq \theta \leq \frac{\pi}{3}$. Thus, we have

$$D^{\mathrm{N}}(u,v) \leq D^{\mathrm{N}}(u,w) + D^{\mathrm{N}}(w,v) \leq c_2 D^{\mathrm{Y}}(u,w) + |wv| \leq c_2 c_1 |uw| + |wv|,$$

where the last step is from the spanner property of the Yao graph. Next, we show that if the min-max distance ratio of λ is λ, then there exists some θ such that $c_1|uw| \leq L(\lambda,\theta)|uv|$, where $L(\lambda,\theta) < 1$ is a constant. We can then choose $c_2 = \frac{1}{1-L(\lambda,\theta)}$. Thus, $D^{\mathrm{N}}(u,v) \leq c_2c_1|uw| + |wv| \leq c_2 L(\lambda,\theta)|uv| + |uv| \leq c_2|uv|$, which will complete the proof of this lemma. Consider triangle $\triangle uvw$. Let $\angle uvw = \alpha$. Given that $\alpha \leq \theta$ and $|wv| \geq \lambda|uv|$, we have $|uw| \leq \max\{2\sin\frac{\theta}{2}, \sqrt{1+\lambda^2-2\lambda\cos\theta}\}|uv|$. Then, one can show that in order for $c_1|uw| = \frac{1}{1-2\sin\frac{\theta}{2}}|uw| < 1$, any θ such that $\sin\frac{\theta}{2} \leq \min\{\frac{1}{4}, \frac{1-\sqrt{1-\lambda(1-\lambda)(2-\lambda)}}{2(1-\lambda)}\}$ suffices. Note that such a angle θ exists for any $0 < \lambda \leq 1$. This completes the proof. ■

Given Lemma 2.6 and spanner property of the Yao graph, Theorem 2 holds.

2.4 Interference model and throughput analysis

Modeling interference in a wireless environment is a complex task. The wireless medium is susceptible to path loss, noise, interference and blockages due to physical obstructions. In this paper, we adopt a *pairwise* interference model (sometimes referred to as the *protocol model* [24]), in which we specify conditions on the distances among participating nodes under which a given transmission is successfully received. Let $X_1, X_2, \ldots, X_k$ be the set of nodes transmitting simultaneously to receivers $Y_1, Y_2, \ldots, Y_k$, respectively, at some instant. Then the transmission by X_i is *successfully received* by node Y_i if $|X_jY_i| \geq (1+\Delta)|X_jY_j|$, for every other node X_j, where $\Delta > 0$ models a protocol specified guard zone to prevent transmission interference.

We consider any message exchange between X_i and Y_i as a bidirectional communication consisting of a transmission from X_i to Y_i and another transmission from Y_i to X_i, to account for both data packets and control packets such as acknowledgments. We define

$$IR(X_i, Y_i) = C(X_i, (1+\Delta)|X_iY_i|) \bigcup C(Y_i, (1+\Delta)|X_iY_i|)$$

to be the *interference region* of transmission $X_i \leftrightarrow Y_i$, where $C(O, r)$ denotes the open disk with center O and radius r. Thus, $X_i \leftrightarrow Y_i$ is successful if and only if for any other transmission $X_j \leftrightarrow Y_j$, both X_i and Y_i are not in $IR(X_j, Y_j)$. We say that an edge e' *interferes* with $e \in E$ if the interference region of e' contains at least one end point of e. Following the recent work of Meyer auf der Heide et al [7], we define the *interference set* of e as $I(e) = \{e' \in E \mid e' \text{ interferes with } e, \text{ or vice versa}\}$, and call $\max_{e \in E}\{I(e)\}$ the *interference number* of the graph.

For an arbitrary communication pattern, the throughput achievable on a given topology depends on both the interference number of the topology and the congestion of the best path system connecting source-destination pairs, both of which in turn are a function of the distribution of the nodes in the plane. In the following theorem, we show the throughput achievable on $\mathcal{N}$ is essentially limited only by its interference number, when compared with an optimal schedule on G^*.

Theorem 3 *Let I be the interference number of topology $\mathcal{N}$. Let W denote a set of packets that are successfully delivered by an arbitrary schedule of packet transmissions in G^* in t steps. Then, there exists a schedule of transmissions in $\mathcal{N}$ that delivers W in $O(tI + n^2)$ steps. Thus, for sufficiently large t and W, the throughput achievable on $\mathcal{N}$ is an $\Omega(1/I)$ fraction of the optimal.*

Let T be any set of edges in G^* such that the any two edges do not interfere with each other. We show in the following that any edge (u, v) in T can be replaced by a set of edges $\{(u, u_1), (u_1, u_2), \ldots, (u_k, v)\}$ from $\mathcal{N}$, such that any edge in $\mathcal{N}$ can be included in at most a constant number of such set of edges in T. We replace an edge (u, v) in G^* by a path P, which is computed as follows. Initially, we have $P = \phi$. If $(u, v) \in E$, then P is simply the edge (u, v). Otherwise, we have two cases. If v is the nearest neighbor of u in $S(u, v)$, then let w be the node in $S(v, u)$ such that (v, w) is an edge. We set P to be the recursive path from u to w, followed by the edge (w, v). If v is not the nearest neighbor of u in $S(u, v)$, then let w be the nearest neighbor of u in $S(v, u)$. We set P to be the recursive path from u to w followed by the recursive path from w to v.

We refer to the path identified by the above algorithm as the θ-path of (u, v). We first establish a key claim that no θ-path diverges, which then implies an upper bound on the overlap among the θ-paths.

Lemma 2.7 *For any $\Delta > 0$, there exists $\theta > 0$ such that the θ-path of (u, v) is bounded by $C(u, (1+\Delta)|uv|)$ on the topology generated by the* ΘALG.

Proof: We prove this lemma by showing that for any intermediate node w selected by an induction step in $Path\_Select(u, v)$, we have the following three properties:

- $w \in C(u, |uv|)$;
- $C(u, (1+\Delta)|uw|) \subset C(u, (1+\Delta)|uv|)$;
- $C(w, (1+\Delta)|wv|) \subset C(u, (1+\Delta)|uv|)$, if $(w, v) \notin E$,

for some $\theta \geq 0$. It is clear that if we have the above properties, the desired lemma follows.

We distinguish between the following two cases:

Case 1: If v is u's nearest neighbor in $S(u, v)$, we select w s.t. $w \in S(v, u)$ and $(w, v) \in E$. Since in this case $(w, v) \in E$, no further induction is needed. The three properties hold.

Case 2: If v is not u's nearest neighbor in $S(u,v)$, there exists w s.t. w is u's nearest neighbor in $S(u,v)$. It is clear that the first and second properties hold. Let $|uv| = 1$, $d = |uw|$ and $\alpha = \angle wuv \le \theta$, thus $|wv| = \sqrt{1+d^2-2\cos\alpha}$. We then show that for any $\Delta > 0$, there exists some $0 < \theta \le \pi/3$, s.t. $C(w,(1+\Delta)|wv|) \subset C(u,(1+\Delta)|uv|)$. Consider the following inequality,

$$d + (1+\Delta)\sqrt{1+d^2-2\cos\theta} \le 1+\Delta.$$

Since $\alpha \le \theta$, and for any $\Delta > 0$ there exists $\theta > 0$, s.t. the above inequality holds. Thus, we have the third property. This completes the proof of the lemma. ■

Lemma 2.8 *Any edge in $\mathcal{N}$ can be selected by at most 6 θ-paths of edges in T.*

Proof: Let o be any point in the 2-dimensional plane. Let $(s_1,r_1),(s_2,r_2)$ be any two edges in T with o in both the interference regions. Consider triangle $\triangle s_1 o s_2$. Since s_1, s_2 do not interfere with each other, we have $|s_1,o| \le |s_1,s_2|$ and $|s_2,o| \le |s_1,s_2|$. Thus, $|s_1,s_2|$ is the largest edge in $\triangle s_1 o s_2$, consequently, $\angle s_1 o s_2$ is at least $\pi/3$. This means that o can be covered by at most 6 interference regions of edges in T. From Lemma 2.7, we know that the θ-path is bounded by its corresponding interference region, thus, any edge in $\mathcal{N}$ can be selected by at most 6 θ-paths. ■

Proof of Theorem 3: Let $\mathcal{S}$ be a transmission schedule on G^* that routes packets in W. In each step of $\mathcal{S}$ a transmission set is specified, which consists of edges in G^* that do not interfere with each other. Consider any step that has a transmission set $T = e_0, e_1, \ldots, e_k$. Let T' be a set of edges in $\mathcal{N}$ by replacing each edge $e_i \in T$ with its θ-path in $\mathcal{N}$.

We construct a schedule $\mathcal{S}'$ for $\mathcal{N}$ by having $6I$ steps with transmission set T' for every step with transmission set T in $\mathcal{S}$. From Lemma 2.8, any edge e in $\mathcal{N}$ is included in at most 6 θ-paths of edges in T; also, e is interfering with at most τ edges in $\mathcal{N}$. Thus, by a simple coloring argument, one can show that there exists $6I$ sets such that each edge is in at least one of these sets. Each node along the θ-path for an edge e has a buffer B_e for storing packets that are transmitted across e in $\mathcal{S}$ but are in transit in $\mathcal{S}'$. We adopt the following convention for forwarding packets: When any edge in the θ-path of edge e is scheduled in $\mathcal{S}'$, the oldest packet in B_e, if any, is forwarded. It can be shown that for $\mathcal{S}'$, the number of packets in the network is at most $O(n^2)$ more than that for $\mathcal{S}$. Thus, $\mathcal{S}'$ can route all of the packets in W within time $O(tI+n^2)$. ■

Finally, we establish an upper bound on the interference number of $\mathcal{N}$ for a random node distribution.

Lemma 2.9 *If the n nodes are placed independently and uniformly at random in a unit square, then the interference number of $\mathcal{N}$ is $O(\log n)$ whp.*

The proof of Lemma 2.9 uses the following claim.

Lemma 2.10 *For any edge $(u,v) \in E$, $|u,v| \le \Theta(\sqrt{\log n/n})$ whp.*

Proof: We consider the directed graph $\mathcal{N}_1 = (V, E_1)$. We show that for any incoming edge e of node u, $|e| \le \Theta(\sqrt{\log n/n})$. Let B denote the circular region $C(u, \sqrt{4c\log n/(\theta n)})$ with center u, where $c > 0$ is a constant. Let random variable Z denote the number of nodes of V in B. It follows that $E[Z] = \pi(\sqrt{4c\log n/(\theta n)})^2 n = 4c\pi\log n/\theta$. We then invoke the following Chernoff bound formula [15]

$$\Pr[Z \le (1-\epsilon)E[Z]] \le e^{-\frac{\epsilon^2}{2}E[Z]}.$$

Let $\epsilon = 1/2$, we have

$$\Pr[Z \le 2c\pi\log n/\theta] \le e^{-\frac{1}{2}c\pi\log n/\theta} = \frac{1}{n^{c\pi/(2\theta)}}.$$

Now consider any circular region with center u containing $2c\pi \log n/\theta$ nodes. Using standard probabilistic analysis, we have that the probability that u can find a neighbor in one sector is at least $1 - 1/n^c$. Consequently, we have

$$\Pr[u \; has \; at \; least \; one \; neighbor \; in \; each \; sector] \geq (1 - \frac{1}{n^c})^{2\pi/\theta} \geq 1 - \frac{2\pi}{\theta n^c}.$$

Note that an incoming edge (v, u) of u is fixed only if v is the nearest neighbor of u in $S(u, v)$. Thus, we have

$$\Pr[B \; contains \; all \; incoming \; edges \; of \; u] \geq (1 - \frac{1}{n^{c\pi/(2\theta)}})(1 - \frac{2\pi}{\theta n^c}) \geq 1 - \frac{1}{n^{c-1}},$$

for sufficiently large n. Computing over all nodes in V, it follows that

$$\Pr[|u, v| \leq \sqrt{4c \log n/(\theta n)}\ , \ \forall (u, v) \in E] \geq (1 - \frac{1}{n^{c-1}})^n \geq 1 - \frac{1}{n^{c-2}}.$$

This completes the proof. ■

Proof of Lemma 2.9: Let (u, v) be any edge in E, and let B denote $C(u, (1+\Delta)\sqrt{4c \log n/(\theta n)}\ ) \cup C(v, (1+\Delta)\sqrt{4c \log n/(\theta n)}\ )$. From the definition of interference set, for any edge $(x, y) \in I(u, v)$, we have

$$\min\{|ux|, |uy|, |vx|, |vy|\} \leq ((1 + \Delta)|uv|).$$

From Lemma 2.10, we have

$$\begin{aligned} E[|I(u,v)|] &\leq E[\text{the number of edges with at least one end in } B] \\ &+ E[\text{the number of } e \notin B \text{ with } |e| \geq \sqrt{4c\log n/(\theta n)}\ ] \\ &\leq 2D\Theta(\log n) + Dn/(2n^{c-2}) = \Theta(\log n), \end{aligned}$$

for sufficiently large n, where D is the maximum degree of N.

We invoke the Chernoff bound [15] as in the proof of Lemma 2.10, we have $|I(u, v)|$ is $O(\log n)$ whp. Computing over all the edges of N, we have that the interference number of N is $O(\log n)$ whp. This completes the proof. ■

3 Routing

In this section we will show how to perform routing in wireless networks to ensure that, in conjunction with certain topology control and medium access control protocols, the throughput and energy efficiency is close to a best possible. After describing our basic model, we investigate various scenarios of comparing optimal algorithms with our algorithms:

1. First, we assume that protocols for topology control and the medium access control are already given. This means that in each step a set of edges is provided that do not interfere with each other and therefore can be used concurrently. MAC layer protocols that allow to achieve this are, for example, CSMA/CA[16], MACA [13, 28] and IEEE 802.11 [26]. Thus, it remains to perform routing decisions to achieve a throughput that is as high as possible.

2. Next, we assume that a topology control protocol is only given, and medium access control and routing protocols have to be designed. In this case, we compare the performance of our algorithm with a best possible routing strategy using the interference number of underlying topology.

3. Finally, we investigate two special cases, one when the ad hoc network nodes are randomly placed in a unit square and the other where the nodes are arbitrarily placed but have a fixed transmission strength and hence have to transmit every packet in the same range.

3.1 Analytical approach

We adopt a model in which the topology changes and packet injections are under adversarial control. That is, in each time step the adversary can specify a new topology with edge costs that may differ from previous edge costs, and it can inject an unbounded number of packets. Of course, in this case only some of the injected packets may be able to reach their destination, even when using a best possible strategy. For each of the successful packets a *schedule* can be specified. A schedule $S = (t_0, (e_1, t_1), \ldots, (e_\ell, t_\ell))$ consists of a sequence of movements by which the injected packet P can be sent from its source node to its destination node. It has the property that P is injected at time step t_0, the edges $e_1, \ldots, e_\ell$ form a connected path, with the starting point of e_1 being the source of P and the endpoint of e_ℓ being the destination of P, the time steps have the ordering $t_0 < t_1 < \ldots < t_\ell$, and edge e_i is active at time t_i for all $1 \le i \le \ell$.

We assume that at most one packet can be transmitted along any edge in each direction and require that no two schedules conflict with each other. That is, no edge is used by two schedules at the same time. When speaking about schedules in the following, we always mean a delivery strategy chosen by a best possible routing algorithm.

We assume that every node v in the system has a buffer $Q_{v,d}$ for each destination d. If a packet reaches its destination buffer $Q_{d,d}$, it is absorbed, and we count it as a successful delivery. The number of deliveries that is achieved by an algorithm is called its *throughput*. Since the adversary is allowed to inject an unbounded number of packets, we will allow routing algorithms to drop packets so that a high throughput can be achieved with a buffer size and a cost that is as small as possible.

In order to compare the performance of an optimal algorithm with our online algorithm, we will use competitive analysis. Given any sequence of topology changes and packet injections σ, let $\mathrm{OPT}_{B,C}(\sigma)$ be the maximum possible throughput (i.e. the maximum number of deliveries) achievable when using a buffer size of B and allowing an average cost of C per delivery (where the average is taken by dividing the total cost spent on all packets by the number of successful deliveries). Let $A_{B',C'}(\sigma)$ be the throughput achieved by some given online algorithm A with buffer size B' and an asymptotic average cost of at most C' per delivery. (Asymptotic means here that as the number of successful deliveries goes to infinity, the average cost goes to at most C'.) We call an online algorithm A (t, s, c)-*competitive* if for all σ and all B and C, A can guarantee that

$$A_{s \cdot B, c \cdot C}(\sigma) \ge t \cdot \mathrm{OPT}_{B,C}(\sigma) - r$$

for some value $r \ge 0$ that is independent of σ (but may depend on s, B and n). Note that $t \in [0, 1]$. For the case that we do not consider the cost of transmissions, we simply say that A is (t, s)-competitive.

3.2 MAC-based routing

We begin with the scenario in which protocols for topology and medium access control are given and it remains to provide a routing protocol. Recall that in this case edges are provided in each step that do not interfere with each other.

The (T, γ)-balancing algorithm

Let $h_{(v,d),t}$ denote the amount of packets in buffer $Q_{v,d}$ at the beginning of time step t. For any destination buffer $Q_{d,d}$, $h_{(d,d),t} = 0$ at any time. $h_{(v,d),t}$ will also be called the *height* of buffer $Q_{v,d}$ at step t. The maximum height a buffer can have is denoted by H.

We now present a simple balancing strategy. Several variants of it have been used in previous papers (e.g. [4, 8, 9]), but without considering edge costs. In every time step $t \ge 1$ the (T, γ)-balancing algorithm performs the following operations.

1. For every edge $e = (v, w)$, determine the destination d with maximum $h_{(v,d),t} - h_{(w,d),t} - c(e) \cdot \gamma$ and check whether $h_{(v,d),t} - h_{(w,d),t} - c(e) \cdot \gamma > T$. If so, send a packet from $Q_{v,d}$ to $Q_{w,d}$ (otherwise do nothing).

2. Receive incoming packets and absorb all packets that reached the destination. Afterwards, receive all newly injected packets. If a packet cannot be stored in a buffer because its height is already H, then delete the new packet.

In the above algorithm, we assume that nodes continuously exchange the buffer height values. In a practical implementation, we can reduce the amount of control information exchange for this purpose. This aspect is discussed in Section 3.5.

Note that if T is set large enough to ensure that packets can only move downwards in their buffer position, then only newly injected packets will ever get deleted. In this case, the admission control problem for the sources has a simple solution: only admit those packets for which there is still buffer space available. We show that this solution is surprisingly effective.

Let $\bar{C}$ denote the average cost allowed for an optimal algorithm to deliver a packet and $\bar{L}$ denote the (best possible) average path length used by successful packets in an optimal algorithm under this assumption. We assume that $\bar{C}$ is known to the online algorithm whereas for $\bar{L}$ just an upper bound may be known. The following result demonstrates that the (T, γ)-balancing algorithm can reach a $(1-\epsilon)$-fraction of the optimal throughput at the cost of increasing the buffer size by a factor of essentially $O(\bar{L}/\epsilon)$ and the average cost per packet by a factor of $O(1/\epsilon)$.

Theorem 3.1 *For any* $\epsilon > 0$ *and any* $T \geq B + 2(\delta - 1)$ *and* $\gamma \geq (T + B + \delta)\bar{L}/\bar{C}$*, the* (T, γ)*-balancing algorithm is* $(1 - \epsilon, 1 + 2(1 + (T + \delta)/B)\bar{L}/\epsilon, 1 + 2/\epsilon)$*-competitive.*

Proof: The proof extends an analysis technique of [9] to incorporate edge costs, and we only present here those parts that differ from a proof given for edges without costs in [9].

Recall that each buffer has H slots to store packets. The slots are numbered in a consecutive way starting from below with 1. Every slot can store at most one packet. After every step of the balancing algorithm we assume that if a buffer holds h packets, then its first h slots are occupied. The *height* of a packet is defined as the number of the slot in which it is currently stored. If a new packet is injected, it will obtain the lowest slot that is available after all packets that are moved to that node from another node have been placed.

Recall that for every successful packet in an optimal algorithm a schedule can be identified. A schedule $S = (t_0, (e_1, t_1), \ldots, (e_\ell, t_\ell))$ is called *active* at time t if $t_0 \leq t \leq t_\ell$. The *position* of a schedule at time t is the buffer at which its corresponding packet would be at that time if it is moved according to S. An edge is called a *schedule edge* if it belongs to a schedule of a packet.

We distinguish between three kinds of packets: *representatives*, *zombies*, and *losers*. During their lifetime, the packets have to fulfill certain rules. (These rules are crucial for our analysis. The balancing algorithm, of course, cannot and does not distinguish between these types of packets.) Every injected packet that does not have a schedule will initially be a zombie. Every other packet will initially be a representative. If a packet is injected into a full buffer and therefore has to be deleted, then the highest available loser will be selected to take over its role.

We want to ensure that a representative always stays with its schedule as long as this is possible. Two cases have to be considered for this when the adversary offers an edge $e = (v, w)$ belonging to a schedule for a packet with destination d:

1. A packet from $Q_{v,d}$ is sent along e: Then we always make sure that this packet is the representative belonging to the schedule.

2. No packet from $Q_{v,d}$ is sent along e: If w has a loser, then the representative exchanges its role with the highest available loser in $Q_{w,d}$. In this case we will also talk about a *virtual* movement. Otherwise, the representative is simply transformed into a loser. In this case, we will disregard the rest of the schedule (i.e. we will not select a representative for it afterwards and the rest of the schedule edges will simply be treated as non-schedule edges).

Furthermore, if a packet from a different buffer $Q_{v,d'}$ is sent along e, then we always make sure that none of the representatives is moved out of $Q_{v,d'}$ but only a loser (which always exists if T is large enough).

The three types of packets are stored in the slots in a particular order. The lowest slots are always occupied by the losers, followed by the zombies and finally the representatives. Every zombie that happens to be placed in a slot of height at most $H - B$ will be immediately converted into a loser. Together with results in [9], these rules allow to prove the following key lemma.

Lemma 3.2 *Let σ be an arbitrary sequence of edge activations and packet injections. Suppose that in an optimal strategy, s of the injected packets have schedules and the other z packets do not. Let$\bar{L}$ be the average length of the schedules. If $H > B$, then the number of packets that are deleted by the balancing algorithm is at most*

$$s \cdot \frac{(T+B+\delta)\bar{L} + \gamma\bar{C}}{H-B} + z \ .$$

Proof. Let $h_{(v,d),t}$ be the height of buffer $Q_{v,d}$ (i.e. the number of packets stored in it) at the beginning of time step t, and let $h'_{(v,d),t}$ be its height when considering only the losers. The *potential* of buffer $Q_{v,d}$ at step t is defined as $\phi_{(v,d),t} = \sum_{j=1}^{h'_{(v,d),t}} j = \binom{h'_{(v,d),t}+1}{2}$ and the potential of the system at step t is defined as $\Phi_t = \sum_{v,d} \phi_{(v,d),t}$.

Let S_t denote the set of all edges belonging to a schedule at time t, $c_t(e)$ be the cost of edge e at time t, and d denote the number of packets that are deleted by the balancing algorithm. Since according to results in [9],

- edges without a schedule do not increase the potential,
- at time t, every edge e belonging to a schedule increases the potential by at most $c_t(e) \cdot \gamma + (T + B + \delta)$,
- every deletion of a newly injected packet decreases the potential by at least $H - B$, and
- every zombie increases the potential by at most $H - B$,

it holds for the potential Φ after executing σ that

$$\Phi \leq \sum_t \sum_{e \in S_t} (c_t(e) \cdot \gamma + T + B + \delta) + z \cdot (H-B) - d \cdot (H-B) \ .$$

Using the fact that for the number of successful deliveries by the optimal scheme, s, it holds that $s = (\sum_t |S_t|)/\bar{L}$ and $\sum_t \sum_{e \in S_t} c_t(e) \cdot \gamma \leq \gamma \cdot s \cdot \bar{C}$, it follows that

$$\Phi \leq s((T+B+\delta)\bar{L} + \gamma \cdot \bar{C}) + z \cdot (H-B) - d \cdot (H-B) \ .$$

Since on the other hand $\Phi \geq 0$, it follows that

$$d \leq s \cdot \frac{(T+B+\delta)\bar{L} + \gamma \cdot \bar{C}}{H-B} + z \ .$$

■

From Lemma 3.2 it follows that the number of packets that are successfully delivered to their destination by the balancing algorithm must be at least

$$s + z - \left(s \cdot \frac{(T+B+\delta)\bar{L} + \gamma \cdot \bar{C}}{H-B} + z \right) - H \cdot N = s \cdot \left(1 - \frac{(T+B+\delta)\bar{L} + \gamma \cdot \bar{C}}{H-B} \right) - H \cdot N \ ,$$

where N is the number of (virtual) non-destination nodes. For $H = ((T+B+\delta)\bar{L} + \gamma \cdot \bar{C})/\epsilon + B$ for some $\epsilon < 1$, this is at least

$$(1-\epsilon)s - r$$

for some value r independent of the number of packets successful in an optimal schedule.

Due to our threshold rule, it holds that when choosing $\gamma \geq (T+B+\delta)\bar{L}/\bar{C}$, every packet in the system can be responsible for creating a cost of at most

$$\frac{H}{\gamma} = \frac{((T+B+\delta)\bar{L} + \gamma \cdot \bar{C})/\epsilon + B}{\gamma} \leq \left(1 + \frac{2}{\epsilon}\right) \bar{C} \ .$$

Since a packet can only get deleted at injection time, every packet that once enters the system either will reach its destination or will get stuck in it. Hence, the total amount of energy spent by the (T,γ)-balancing algorithm is at most

$$(s' + H \cdot N) \cdot \left(1 + \frac{2}{\epsilon}\right) \bar{C}$$

where s' is the number of successful deliveries of our algorithm. Hence, the average cost per delivered packet is equal to

$$\frac{s' + H \cdot N}{s'} \cdot \left(1 + \frac{2}{\epsilon}\right) \bar{C} \overset{s \to \infty}{\to} \left(1 + \frac{2}{\epsilon}\right) \bar{C} \ .$$

■

3.3 Topology-based routing

Next we show that it is possible to compete with an optimal algorithm even when medium access control is not provided. Recall the definition of the interference number in Section 2.4. Suppose that we use a topology control algorithm such as ΘALG of Section 2, and suppose that every node v knows for every edge $e = (v,w)$ of the resultant topology an upper bound I_e on the maximum current interference number of any edge e interferes with. (In the ideal, 2-dimensional Euclidean space it would actually suffice just to have an upper bound on the own interference number, but in a space with obstacles, for example, this would not suffice.) Then we use the following simple symmetry-breaking technique to provide medium access control: Each edge e provided by the topology control scheme chooses to become active with probability $1/(2I_e)$. All active edges are passed on to the (T,γ)-balancing algorithm. We refer to the combined medium access and routing protocol as a (T,γ,I)-*balancing* algorithm, where I is the maximum of I_e over all edges every offered by the topology control protocol.

We note that if the algorithm decides to send packets along two active edges that interfere with each other, then neither of transmissions is successful. Fortunately, the following lemma demonstrates that there is a high probability of successfully using an active edge.

Lemma 3.3 *Every active edge has a probability of at most* $1/2$ *to interfere with other active edges.*

Proof. Consider any such edge e, and let N_e be its current interference number. Since for all edges e' it interferes with, $I_{e'} \geq N_e$ according to the definition of $I_{e'}$, it follows that the probability that at least one of these edges is active is at most $N_e \cdot 1/(2N_e) = 1/2$. ■

Now we are ready to compare our (T,γ,I)-balancing algorithm with the performance of an optimal algorithm. We assume here that an optimal algorithm still has to restrict itself to the edges provided by the topology control scheme, but apart from that is free to use any of these edges for communication as it likes. We even allow it to use edges successfully at the same time that would normally interfere with each other. Let δ, the maximum degree of a node in a step, be now equal to 1 (i.e. only one frequency is available).

Theorem 3.4 *For any* $\epsilon > 0$ *and any* $T \geq 2B+1$ *and* $\gamma \geq (T+B)\bar{L}/\bar{C}$*, the* (T,γ,I)*-balancing algorithm is* $((1-\epsilon)/(8I), 1+2(1+T/B)\bar{L}/\epsilon, 1+2/\epsilon)$*-competitive w.r.t. an optimal algorithm that is based on the same topology control scheme.*

Proof: In order to compare the (T, γ, I)-balancing algorithm with an optimal algorithm, OPT, we use a trick to compare it with a randomized algorithm OPT' derived from OPT in a way that every schedule in OPT is only used by OPT' with probability $1/(8I)$.

Recall the notation used in the proof of Theorem 3.1. In contrast to that proof, we will use the rule that for every active edge that can be used without interference, we always move a loser, and only virtually move representatives if necessary and possible.

Consider some fixed step t. Given an edge $e = (v, w)$, let its *benefit* at time t be defined as $b_e = \max_d(h_{(v,d),t} - h_{(w,d),t})$. Let B_t be the set of all edges provided by the topology control scheme at time t with $b_e > T + c(e) \cdot \gamma$, and let O_t be the set of all edges used by schedules in OPT at time t. Furthermore, let O'_t consist of all edges $e \in O_t$ with $b_e > T + c(e) \cdot \gamma$ and let $O''_t = O_t \setminus O'_t$. First we consider the potential decrease caused by edges in B_t.

Lemma 3.5 *The expected potential decrease caused by B_t is at least* $\frac{1}{4I} \sum_{e \in B_t} (b_e - (B + 1))$.

Proof. Consider any edge $e = (v, w) \in B_t$. Since $b_e > T + \gamma \cdot c(e)$, the (T, γ, I)-balancing algorithm would send a packet along e if successfully activated, using any buffer pair with height difference b_e. Let d be the destination associated with that buffer pair. Since we have to make sure not to move a representative away from its position, we have to take the highest loser in $Q_{v,d}$ and move it to $Q_{w,d}$. This causes a potential decrease of at least

$$(h_{(v,d),t} - B) - (h_{(w,d),t} + 1) = (h_{(v,d),t} - h_{(w,d),t}) - (B + 1) = b_e - (B + 1) \ .$$

Since any edge in B_t is successfully used with a probability of at least $1/2 \cdot 1/(2I) = 1/(4I)$, the lemma follows. ■

The next lemma gives a bound on the potential increase caused by O''_t.

Lemma 3.6 *The potential increase caused by O''_t is at most* $\sum_{e \in O''_t} (c(e) \cdot \gamma + T + B)$.

Proof. Note that the edges in O''_t are not considered by the balancing algorithm and therefore $B_t \cap O''_t = \emptyset$. Consider some edge $e = (v, w) \in O''_t$ and let R be its representative and d be the destination of R. (If there is no such R because the representative belonging to the schedule owning e has already been discarded, we do not have to perform a virtual movement, and therefore the potential will not increase due to e.) Let h_R be the height of R in $Q_{v,d}$ and let h_L be the height of the highest loser in $Q_{w,d}$. If there is no such loser, then $Q_{w,d} \leq B$ and therefore $Q_{v,d} \leq c(e) \cdot \gamma + T + B$ (because $b_e \leq T + \gamma \cdot c(e)$). In this case, we would simply transform R into a loser, causing a potential increase by at most $c(e) \cdot \gamma + T + B$. Otherwise, we replace the roles of R and the loser L. Since w can lose at most one packet during step t, this causes a potential increase of at most

$$h_{(v,d),t} - (h_{(w,d),t} - B) \leq b_e + B \leq c(e) \cdot \gamma + T + B \ ,$$

which completes the lemma. ■

It remains to bound the potential increase caused by the edges in O'_t.

Lemma 3.7 *The potential increase caused by the edges in O'_t is at most* $\sum_{e \in O'_t} (c(e) \cdot \gamma + b_e + B)$.

Proof. Consider any edge $e = (v, w) \in O'_t$, and let R be the representative OPT would move from $Q_{v,d}$ to $Q_{w,d}$ for some d. Recall that we want to move all representatives belonging to the schedule edge in O'_t virtually. Since $Q_{w,d}$ can lose at most one packet in step t, the potential increase by virtually moving R is at most

$$h_{(v,d),t} - (h_{(w,d),t} - B) = (h_{(v,d),t} - h_{(w,d),t}) + B \leq c(e) \cdot \gamma + b_e + B \ .$$

■

Now we are ready to prove an upper bound on the number of packets that are deleted by the (T, γ, I)-balancing algorithm.

Lemma 3.8 *Let σ be an arbitrary sequence of topology selections and packet injections. Suppose that in* OPT*, s of the injected packets have schedules and the other z packets do not. Let $\bar{L}$ be the average length of the schedules. If $T \geq 2B+1$ and $H > B$, then the expected number of packets that are deleted by the (T, γ, I)-balancing algorithm is at most*

$$\left(1 - \frac{1}{8I} + \frac{(T+B)\bar{L} + \gamma\bar{C}}{8I(H-B)}\right) s + z \ .$$

Proof. First of all, note that only newly injected packets get deleted. Let p denote the number of schedule edges in OPT and d denote the number of packets that are deleted by the balancing algorithm.

If $T \geq 2B$, then it follows from Lemma 3.5 and the fact that $\sum_{e \in B_t} b_e \geq \sum_{e \in O'_t} b_e$ that the expected potential decrease at step t is at least

$$\frac{1}{4I} \sum_{e \in B_t} (b_e - (B+1)) \geq \frac{1}{8I} \sum_{e \in B_t} b_e \geq \frac{1}{8I} \sum_{e \in O'_t} b_e \ . \tag{1}$$

On the other hand, it follows from Lemma 3.7 that the expected potential increase at step t due to OPT' is at most

$$\frac{1}{8I} \left(\sum_{e \in O'_t} (c(e) \cdot \gamma + b_e + B) + \sum_{e \in O''_t} (c(e) \cdot \gamma + T + B) \right) \ . \tag{2}$$

Combining (1) and (2), the expected potential increase w.r.t. schedule edges used by OPT' at step t is at most

$$\frac{1}{8I} \left(\sum_{e \in O'_t} (c(e) \cdot \gamma + b_e + B) + \sum_{e \in O''_t} (c(e) \cdot \gamma + B + T) - \sum_{e \in O'_t} b_e \right) \leq \frac{1}{8I} \sum_{e \in O_t} (c(e) \cdot \gamma + B + T) \ .$$

Since $\sum_t \sum_{e \in O_t} c_t(e) \cdot \gamma \leq \gamma \cdot s \cdot \bar{C}$ and $\sum_t |O_t| = p$, it follows that the overall expected potential increase w.r.t. schedule edges used by OPT' is at most

$$\frac{1}{8I} \left(\gamma \cdot s \cdot \bar{C} + p(B+T) \right) \ .$$

Furthermore, it follows from results in [9]

- every deletion of a newly injected packet decreases the potential by at least $H - B$, and
- every zombie increases the potential by at most $H - B$.

Hence, it holds for the expected value of the potential Φ after executing σ that

$$\mathrm{E}[\Phi] \leq \frac{1}{8I} \left(\gamma \cdot s \cdot \bar{C} + p(B+T) \right) + \mathrm{E}[z'] \cdot (H-B) - \mathrm{E}[d] \cdot (H-B)$$

where z' is the number of zombies in OPT'. Since on the other hand $\mathrm{E}[\Phi] \geq 0$, it follows that

$$\mathrm{E}[d] \leq \frac{\gamma \cdot s \cdot \bar{C} + p \cdot (T+B)}{8I(H-B)} + \mathrm{E}[z'] \ .$$

Using in this inequality the facts that $\mathrm{E}[z'] = (1 - 1/(8I))s + z$ and that $p = s \cdot \bar{L}$, the lemma follows. ∎

From Lemma 3.8 it follows that the expected number of packets that are successfully delivered to their destination by the (T, γ, I)-balancing algorithm must be at least

$$s + z - \left(\left(1 - \frac{1}{8I} + \frac{(T+B)\bar{L} + \gamma\bar{C}}{8I(H-B)} \right) s + z \right) - H \cdot N = \frac{s}{8I} \cdot \left(1 - \frac{(T+B)\bar{L} + \gamma\bar{C}}{H-B} \right) - H \cdot N \ ,$$

where N is the number of buffers in the network. For $H = ((T+B)\bar{L} + \gamma\bar{C})/\epsilon + B$ for some $\epsilon < 1$ this is at least

$$\frac{1-\epsilon}{8I} \cdot s - r$$

for some value r independent of the number of packets successful in an optimal schedule. The bound on the energy efficiency can be shown in the same way as for Theorem 3.1. ∎

Theorem 3.4 can be combined with an analysis along the lines of Theorem 3 to yield the following claim for ΘALG and the (T, γ, I)-balancing algorithm, when compared with an optimal algorithm that is unrestricted in what edges it can use.

Corollary 3.9 *Suppose the nodes in the ad hoc network are stationary and the adversary only controls packet injections. For suitable values of T and γ, the (T, γ, I)-balancing algorithm, in conjunction with ΘALG, is $(O(1/I), O(\bar{L}))$-competitive w.r.t. an optimal algorithm that may use any edges of G^*.*

For the special case of having a random distribution of nodes in the unit square, Corollary 3.9 and Lemma 2.9 imply the following:

Corollary 3.10 *Suppose the nodes in the ad hoc network are randomly distributed in a unit square and the adversary only controls packet injections. For suitable values of T and γ, the (T, γ, I)-balancing algorithm, in conjunction with ΘALG, is $(O(1/\log n), O(\bar{L}))$-competitive w.r.t an optimal algorithm that may use any edges of G^*.*

3.4 Communication with fixed transmission strength

Finally, we demonstrate that an even better competitive ratio than the one given in Corollary 3.10 can be shown if the nodes are distributed in an arbitrary way in a 2-dimensional Euclidean space but all nodes have the same, fixed transmission strength. That is, we assume that every node transmits at the same fixed power level so that it will be successfully received by all nodes within distance 1, if there were no interference. Now recall the interference model in Section 2.4. For a transmission from a sender s to a receiver t to be successful, two properties have to be kept: (i) the distance between s and t is at most 1, and (ii) every node in every other sender-receiver pair must have a distance of more than $1 + \Delta$ from s and t. If (ii) holds for two sender-receiver pairs, they are said to be *independent*.

Consider now the 2-dimensional space to be partitioned into a honeycomb-like hexagonal pattern as shown in Figure 9, with hexagons of side length $3 + 2\Delta$ (and therefore diameter $2(3 + 2\Delta)$). Each sender-receiver pair (s, t) is assigned to that hexagon that contains s. Our strategy for selecting independent sender-receiver pairs is rather simple: Suppose that every sender-receiver pair has a benefit associated with it, which equals the maximum difference in buffer heights, over all destination buffers. Within each hexagon, we first determine the sender-receiver pair of maximum benefit (breaking ties in an arbitrary way). If this sender-receiver pair has a benefit of more than some threshold $T > 0$, it is called a *contestant*. For each contestant, we decide with probability p_t to transmit a packet along its connection, where p_t is chosen so that the probability of a successful transmission is at least $1/2$. Two important lemmata can be shown for this strategy.

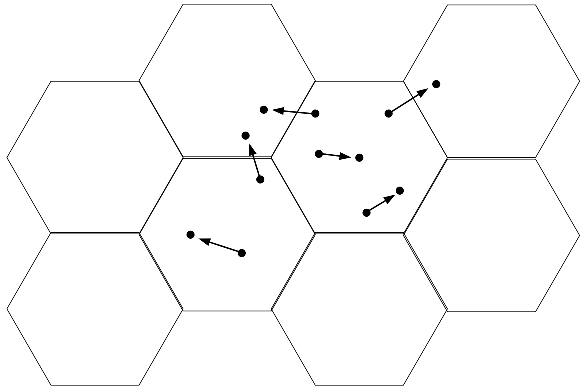

Figure 9: Subdivision of the Euclidean space into hexagons.

Lemma 3.11 *The sum of the benefits of all contestants is by at most some constant factor c_0 smaller than the maximum total benefit that can be achieved by any independent set of sender-receiver pairs with benefit beyond* T.

Proof: Consider some fixed hexagon H and some fixed sender-receiver pair (s,t) in H. Let the *forbidden zone* of (s,t) be defined as the area covered by the disks of radius $\frac{1}{2}(1+\Delta)$ around s and t. Certainly, for every two independent pairs (s,t) and (s',t') in H, their forbidden zones are not allowed to overlap. Since the forbidden zone of every pair belonging to H covers an area or $\Theta(1+\Delta)$ in H and altogether H only covers an area of $\Theta(1+\Delta)$, the largest possible independent set in H can have at most some constant number c_0 of pairs. Since the contestant for H has the maximum benefit of a pair in H, the total benefit of any independent set in H with individual benefits beyond T can be by at most a factor of c_0 larger than the benefit of the contestant. Since the sum over all hexagons H of the maximum total benefit achievable by independent pairs of benefit beyond T in H is at least the maximum total benefit that can be achieved by any globally independent set of pairs with benefit beyond T, the lemma follows. ■

Lemma 3.12 *If $p_t \leq 1/6$, then for each contestant (s,t), with probability at least $1/2$ no other contestant is selected that interferes with (s,t).*

Proof: Consider some contestant (s,t) in some hexagon H. Since every hexagon has a side length of $3+2\Delta$, (s,t) can only interfere with contestants in neighboring hexagons, because s must have a distance of at least $3+2\Delta$ with the source of any contestant in non-neighboring hexagons. So consider any neighboring hexagon H' of H, and let H'' be the neighbor of H on the opposite side of H'. If (s,t) wants to interfere with contestants in H' and H'', it must have a distance of at most $2+\Delta$ to H' and H''. However, since H' and H'' have a distance of $\sqrt{(6+4\Delta)^2-(3+2\Delta)^2} > 5+3\Delta$ and s and t can have a distance of at most 1, this is not possible. This immediately implies that (s,t) can interfere with contestants in at most three neighboring hexagons, because otherwise there would have to be two of them that are on opposite sides of H. Thus, if $p_t \leq 1/6$, then the probability that at least one of the at most three interfering contestants is selected for a transmission is at most $3/6 = 1/2$, which proves the lemma. ■

Let the *honeycomb algorithm* be a combination of the contestant selection strategy and the $(T,\gamma,3)$-balancing algorithm applied to the contestants. The two lemmata above and Theorem 3.4 then yield the following result.

Theorem 3.13 *For any $\epsilon > 0$ and any $T \geq 2B+1$, the honeycomb algorithm is $((1-\epsilon)/(24\,c_0), 1+(1+T/B)L/\epsilon, 1+2/\epsilon)$-competitive.*

3.5 A note on the implementation of the (T,γ)- and (T,γ,I)-balancing algorithms

Finally, we discuss an implementation detail concerning the local balancing algorithms presented in Sections 3.2, 3.3, and 3.4. This is concerning the benefit of an incident edge $e=(v,w)$, resp. the destination d with maximum $h_{(v,d),t}-h_{(w,d),t}-c(e)\cdot\gamma$. If there are sufficiently few destinations in the network and sufficiently large packets, then the amount of data necessary to transmit all of the current values for the $h_{w,t}$'s can be assumed to be just a fraction of the data that has to be transmitted for a packet. In this case, it can be shown that instead of achieving a $1-\epsilon$ throughput, still a constant fraction of the optimal throughput can be achieved (the overhead of sending the $h_{(w,d),t}$ values has a similar effect on the throughput as the interference number).

Another approach to solve the problem to determine the benefit is to update information about the $h_{(w,d),t}$ values only once in a while and therefore to work with outdated information. It is not difficult to show that this does not matter as long as T is chosen sufficiently large so that it is guaranteed that the actual value of $h_{(v,d),t}-h_{(w,d),t}-c(e)\cdot\gamma$ is always more than $2B+1$ if the difference when using the outdated value for $h_{(w,d),t}$ is more than T. Such an approach has, for instance, been successfully used by Aiello et al. [4].

Finally, another option would be to require w to send all $h_{(w,d),t}$ values to v when the connection (v,w) is established. In this case, it would afterwards suffice to transmit only the *changes* that occurred to the $h_{w,d}$ values in w to v. Assuming that w can only send out a single packet per time step, the changes would be sufficiently small so that the amount of data necessary to notify v about them can be regarded as negligible.

References

[1] M. Adler and C. Scheideler. Efficient communication strategies for ad-hoc wireless networks. *Theory of Computing Systems*, 33:337–391, 2000.

[2] Y. Afek and E. Gafni. End-to-end communication in unreliable networks. In *Proc. of the 7th IEEE Symp. on Principles of Distributed Computing (PODC)*, pages 131–148, 1988.

[3] W. Aiello, B. Awerbuch, B. Maggs, and S. Rao. Approximate load balancing on dynamic and synchronous networks. In *Proc. of the 25th ACM Symp. on Theory of Computing (STOC)*, pages 632–641, 1993.

[4] W. Aiello, E. Kushilevitz, R. Ostrovsky, and A. Rosén. Adaptive packet routing for bursty adversarial traffic. In *Proc. of the 30th ACM Symp. on Theory of Computing (STOC)*, pages 359–368, 1998.

[5] M. Andrews, B. Awerbuch, A. Fernández, J. Kleinberg, T. Leighton, and Z. Liu. Universal stability results for greedy contention-resolution protocols. In *Proc. of the 37th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 380–389, 1996.

[6] S. Arya and M. Smid. Efficient construction of a bounded degree spanner with low weight. In *Proceedings of 2nd European Symposium on Algorithms*, pages 48–59, 1994.

[7] F. Meyer auf der Heide, C. Schindelhauer, K. Volbert, and M. Grunewald. Energy, congestion and dilation in radio networks. In *Proceedings of the 14th Annual ACM Symposium on Parallel Algorithms and Architectures*, pages 230–237, August 2002.

[8] B. Awerbuch, P. Berenbrink, A. Brinkmann, and C. Scheideler. Simple routing strategies for adversarial systems. In *Proc. of the 42nd IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 158–167, 2001.

[9] B. Awerbuch, A. Brinkmann, and C. Scheideler. Anycasting and multicasting in adversarial systems. Technical report, Johns Hopkins University, 2002. An electronic copy is available from http://www.cs.jhu.edu/∼scheideler.

[10] B. Awerbuch and F.T. Leighton. A simple local-control approximation algorithm for multicommodity flow. In *Proc. of the 34th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 459–468, 1993.

[11] B. Awerbuch and F.T. Leighton. Improved approximation algorithms for the multi-commodity flow problem and local competitive routing in dynamic networks. In *Proc. of the 26th ACM Symp. on Theory of Computing (STOC)*, pages 487–496, 1994.

[12] B. Awerbuch, Y. Mansour, and N. Shavit. End-to-end communication with polynomial overhead. In *Proc. of the 30th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 358–363, 1989.

[13] V. Bharghavan, A. Demers, S. Shenker, and L. Zhang. MACAW: A media access protocol for wireless LANs. In *Proc. of the 1994 SIGCOMM Conference*, pages 212–225, 1994.

[14] A. Borodin, J. Kleinberg, P. Raghavan, M. Sudan, and D. P. Williamson. Adversarial queueing theory. In *Proc. of the 28th ACM Symp. on Theory of Computing (STOC)*, pages 376–385, 1996.

[15] H. Chernoff. A measure of the asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Annals of Mathematical Statistics*, 23:493–509, 1952.

[16] A. Colvin. CSMA with collision avoidance. *Computer Communications*, 6(5):227–235, 1983.

[17] D. Eppstein. Spanning trees and spanners. In J. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 425–461. Elsevier, Amsterdam, 2000.

[18] J. Freebersyser and B. Leiner. A DoD perspective on mobile ad hoc networks. In C. Perkins, editor, *Ad Hoc Networking*, pages 29–51. Addison-Wesley, 2001.

[19] D. Gamarnik. Stability of adversarial queues via fluid models. In *Proc. of the 29th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 60–70, 1998.

[20] D. Gamarnik. Stability of adaptive and non-adaptive packet routing policies in adversarial queueing networks. In *Proc. of the 31st ACM Symp. on Theory of Computing (STOC)*, pages 206–214, 1999.

[21] J. Gao, L. Guibas, J. Hershberger, L. Zhang, and A. Zhu. Geometric spanner for routing in mobile networks. In *Proceedings of the ACM Symposium on Mobile Ad Hoc Networking & Computing*, October 2001.

[22] A. Goel. Stability of networks and protocols in the adversarial queueing model for packet routing. In *Proc. of the 10th ACM Symp. on Discrete Algorithms (SODA)*, pages 911–912, 1999.

[23] M. Grünewald, T. Lukovszki, C. Schindelhauer, and K. Volbert. Distributed maintenance of resource efficient wireless network topologies. *Draft paper (extended abstract)*, 2002.

[24] P. Gupta and P. Kumar. Capacity of wireless networks. *IEEE Transactions on Information Theory*, IT-46:388–404, 2000.

[25] Y. Hassin and D. Peleg. Sparse communication networks and efficient routing in the plane. In *Proceedings of the 19th Annual ACM Symposium on Principles of Distributed Computing*, pages 41–59, July 2000.

[26] IEEE. Medium access control (mac) and physical specifications. *IEEE P802.11/D10*, January 1999.

[27] H.B. Hunt III, M.V. Marathe, V. Radhakrishnan, S.S. Ravi, D.J. Rosenkrantz, and R. S tearns. NC-approximation schemes for NP- and PSPACE-hard problems for geometric graphs. *Journal of Algorithms*, 26:238–274, 1999.

[28] P. Karn. MACA – A new channel access method for packet radio. In *ARRL/CRRL Amateur Radio 9th Computer Networking Conference*, pages 134–140, 1990.

[29] B. Karp and H. T. Kung. GPSR: greedy perimeter stateless routing for wireless networks. In *Proceedings of ACM Symposium on Mobile Computing and Networking*, pages 243–254, August 2000.

[30] J. M. Keil and C. A. Gutwin. Classes of graphs which approximate the complete euclidean graph. *Discrete and Computational Geometry*, 7:13–28, August 1992.

[31] L. Li, J. Halpern, V. Bahl, Y.-M. Wang, and R. Wattenhofer. Analysis of a cone-based distributed topology control algorithms for wireless multi-hop networks. In *Proceedings of ACM Symposium on Principles of Distributed Computing*, pages 264–273, August 2001.

[32] X.Y. Li, P.J. Wan, Y. Wang, and O. Frieder. Sparse power efficient topology for wireless networks. *Journal of Parallel and Distributed Computing*, 2002.

[33] C. Perkins. Ad hoc networking: An introduction. In C. Perkins, editor, *Ad Hoc Networking*, pages 1–28. Addison-Wesley, 2001.

[34] R. Ramanathan and R. Rosales-Hain. Topology control of multihop wireless networks using transmit power adjustment. In *Proceedings of IEEE Infocom 2000*, March 2000.

[35] T. Rappaport. *Wireless Communications: Principles and Practices*. Prentice Hall, 1996.

[36] J. Salowe. Euclidean spanner graphs with degree four. *Discrete Applied Mathematics*, 54:55–66, 1994.

[37] C. Scheideler and B. Vöcking. From static to dynamic routing: Efficient transformations of store-and-forward protocols. In *Proc. of the 31st ACM Symp. on Theory of Computing (STOC)*, pages 215–224, 1999.

[38] M. Steenstrup. Cluster-based networks. In C. Perkins, editor, *Ad Hoc Networking*, pages 75–138. Addison-Wesley, 2001.

[39] G. Toussaint. The relative neighborhood graph of a finite planar set. *Pattern Recognition*, 12:261–268, 1980.

[40] P. Tsaparas. Stability in adversarial queueing theory. Master's thesis, Dept. of Computer Science, University of Toronto, 1997.

[41] P. Varaiya and J. Walrand. *High-Performance Communication Networks*. Morgan Kaufmann, 2000. Chapter 7.

[42] Y. Wang and X.Y. Li. Distributed spanner with bounded degree for wireless ad hoc networks. In *Parallel and Distributed Computing Issues in Wireless networks and Mobile Computing*, April 2002.

[43] R. Wattenhofer, L. Li, P. Bahl, and Y.M. Wang. Distributed topology control for power efficient operation in multihop wireless ad hoc networks. In *Proceedings of IEEE Infocom*, 2001.

[44] A. Yao. On constructing minimum spanning trees in k-dimensional spaces and related problems. *SIAM Journal on Computing*, 11(4):721–736, 1982.